

ABSTRACT

Large messages in the form of hierarchically structured documents are processed in a streaming fashion using the ultimate consumer read requests as the driving force for the processing. The messages are partitioned into fixed length segments. The segments are processed in pipeline fashion. This processing chain includes simulating random access of hierarchical documents using stream transformations, mapping streams to a transport's native capabilities, composing streams into chains and using pipeline processing on the chains, staging fragments into a database and routing messages when complete messages have been formed, and providing tools to allow the end user to inspect partial messages.